

## DIGITAL MEDIA

A Seybold Report

## Inside

#### 21/0

Bob Metcalfe, the inventor of Ethernet, looks at 3DTV and wonders if networks will be the bottleneck in digital media.

## 12 Apple's new consumer strategy sounds familiar

Through expertise and partnerships, Apple tries to bridge the computer and consumer electronics worlds.

## 14 Winter CES is short on thrills . . .

... but digital technology is poised for a big splash with Apple's new PDAs, DCC, SkyPix and more.

#### 15 Digital Media moves office

## 17 FCC allocates spectrum to interactive

Commission changes the rules of the game for local applicants.

#### 17 'TV Answer' begs the question

Is it a case of market focus or limited vision?

## **19 Correction** ICOM not in video alliance.

### 20 Mediascape

This month, we detail video compression — from MPEG to DVI to QuickTime.

#### 24 Events

## Electronic Publishing on the 'Net' 3

Today, vast stores of information are freely available through a worldwide series of interconnected digital networks called the Internet. The Internet paradigm is increasingly being eyed by information publishers, telephone companies and computer vendors as a powerful platform for a new form of electronic publishing, where the network is more than a means to move data from the desktop to a printer. Instead, the network is where the information resides.

But what happens when information — the contents of a public library, for example, most of which is protected by copyright laws designed for print — is digitized and posted on public networks? At this point, not much. And what we do about electronic copyrights today will have great effect on the electronic publishing market of the future.

## WhereThere's a Will, There's a WAIS 5

In 1989, Thinking Machines Corp. — the company that builds the Connection Machine supercomputer — set Brewster Kahle loose on the problem of catalyzing a market for the electronic distribution of information.

First and foremost, the infrastructure had to allow people to make money at electronic publishing. Second, it needed to be elastic enough to support anything from personal computers to consumer electronics devices to supercomputers. And third, it needed to be completely accessible and autonomous — *i.e.*, there was to be no single point of control.

Kahle's project, dubbed WAIS (pronounced "ways"), for Wide Area Information Servers, is already well on its way to achieving those goals.

## The Reading Room: A Model for Access 7

In this provocative essay, WAIS project director Brewster Kahle offers his views on the clash between private ownership of information, the foundation of the so-called Information Age, and public access to information, which has been the cornerstone of democracy. He suggests we investigate a new model for public libraries, called the Reading Room.

## DIGITAL MEDIA

Publisher, Jonathan Seybold (Malibu)
Editor, Denise Caruso (San Francisco)
Editorial Assistant, Amy Johns
Associate Editor, Peter Dyson (Media)
Associate Editor, David Baron (Malibu)
Contributing Editor, Bernard Banet
Contributing Editor, Connie Guglielmo
Production, John Butler, Preecha Edwards

#### **Business and Production Office**

428 E. Baltimore Pike PO Box 644 Media, PA 19063, USA Phone: (215) 565-2480 Fax: (215) 565-4659, 565-3261

#### San Francisco Office

444 De Haro St., Suite 128 San Francisco, CA 94107, USA Phone: (415) 575-3775 Fax: (415) 575-3780

#### Malibu Office

29160 Heathercliff Road, Suite 200 PO Box 578 Malibu, CA 90265-0578, USA Phone: (310) 457-8500 Fax: (310) 457-8599

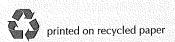
## Overseas Subscriptions UK Office

26A Seaside Eastbourne, East Sussex, BN22 7QJ Phone: {44} 323 410561 Fax: {44} 323 410279

#### **Tokyo Office**

High Technology Communications, Inc. Hachiko Bldg., 1-34-6 Takadanobaba, Shinjuku-ku, Tokyo 160, Japan Phone: 03-3232-7391 Fax: 03-3232-7395







#### READERS RESPOND

## Will networks be the bottleneck in digital media?

## After confronting 3DTV, Ethernet's inventor thinks we should be worried

Bob Metcalfe invented the Ethernet computer networking scheme at Xerox in 1973 and founded 3Com Corp. in 1979. He is now a visiting fellow in the Computer Laboratory at the University of Cambridge, England.

We all learned long ago from personal computers just how interesting interactive software can be without networks, working mostly on small personal databases, and sometimes on large static ones. And we learned recently just how much *more* interesting PCs are when they are networked, augmenting interpersonal communication and giving shared access to large dynamic databases.

Delivering media. Now consider the sum of accelerating developments in digital media — video telephones, HDTV, computer graphics, virtual reality, digital imaging, ubiquitous computing, interactive books, wireless palmtops, video conferencing, electronic publishing, dynamic multimedia, 3DTV — and ponder how these real-time and data-intensive technologies might be delivered. Dig-

ital media without networks may not be media at all.

You probably agree, but have you thought enough about the networks needed to support new media? As a networking specialist, I worry that unless you think more about networks now, they will turn out to be your bottlenecks later, just when digital media approach commercial reality.

Calling the question. I want to get you thinking about what kinds of networks digital media will need. Who will make the huge investments, under what standards, and when, to establish a networking infrastructure? Are current networks as much of a bottleneck as they seem? What applications will drive network development?

I will use my favorite computer network, Ethernet, and a proposed benchmark application, 3DTV—not to answer whether networks will be digital media bottlenecks, but to leave the question churning in your mind.

## Appropriate technology

By "networks" do we mean broadcast television, cable TV, telephones, computer networks? Well, yes, all of them, especially as they are converging. Certainly all are going digital, but is this convergence enough?

Given the magnitude of investments in network infrastructure, and given the concern for standards that such investments imply, the convergence will be slow and, if you do not get involved today, the convergence may be toward networks inappropriate for digital media.

I am a computer networker, and so it is no surprise that I think computer networks are the closest of (continued on page 9)





## Electronic publishing on 'The Net'

## Changing copyright law may help spur the market

Today, vast stores of information are freely available through a worldwide series of interconnected digital networks called the Internet.

The Internet paradigm is increasingly being eyed by information publishers, telephone companies and computer vendors as a powerful platform for a new form of electronic publishing, where the network is more than a means to move data from the desktop to a printer.

Instead, it is where the information resides.

**Sellers beware**. But what happens when information that once was distributed freely — the contents of a public library, for example, most of which is "protected" by copyright laws designed for print — is digitized and posted on public networks?

Will the ability of digital technology to quickly and easily search out (and copy and transmit) information online, at home or office, make publishers impose a per use price structure on information that once had a fixed price?

### Copyright: the big question

The best and simplest answer for the time being is to build a network architecture that can find and retrieve both free and "for pay" information. (See WAIS story, page 5.) But simply figuring out how to ring up the cash register is putting the cart before the horse when it comes to online publishing.

Though traditional publishers have been eyeing the net as as a potential distribution medium for many years, they still greatly fear losing control of their property on a digital network, where making perfect copies or altering existing material is as easy as pressing a key. (See related story on the digital audio industry, page 15.)

It's clear that today's copyright and intellectual property laws, especially for literary works, don't properly serve an information economy, but no one knows exactly what to replace them with. The choices are to sit back and wait until someone else figures it out or jump in and get muddied up yourself.

### Project Gutenberg jumps right in

Project Gutenberg is not afraid of a little mud. The goal of the project, started by an Urbana, IL systems analyst named Michael Hart, is to give away a trillion books — 10,000 "etexts," or electronic books, to 100 million people — by the end of 2001. Its very existence throws down the gauntlet to the publishing system as it exists today.

In fact, all of the project's formal electronic mail messages sign off with the following quotation from Grolier Electronic Publishing: "The trend of library policy is clearly toward the ideal of making all information available without delay to all people."

The goal of the project is to give away a trillion books to 100 million people by the end of 2001.

These words are not music to the ears of the publishing business. Neither would be this quote from the Project's December newsletter: "A great advantage of the computer networks is that we only post a book once and a million people can make their own copies very easily and quickly."

**Camp Free.** Obviously Hart is of the "free information" camp, and his intention is to reduce the effective cost of an etext to about one cent. An entire 10,000-book electronic library would thus cost about \$100 (plus the price of diskettes or CD-ROMs and mailing).

To date, there are more than 25 etexts in Project Gutenberg's "stacks," including versions of *Alice in Wonderland*, *Peter Pan* (which can only be distributed in the U.S. because of copyright restrictions), *The Book of Mormon*, *Far From the Madding Crowd*, *Aesop's Fables* and various reference books and a special etext edition of William Shakespeare's works called *Shakespeare of Disk*.

## Move more quickly to public domain

Hart and others who work with him are adamant that the publishing system as it exists today is a millstone around the neck of progress for electronic publishing.

One of the Project's volunteers, Mary Brandt Jensen, is an associate professor of law and the director of the Law Library at the University of South Dakota. She's taken on the task of clearing copyrights for Gutenberg's etexts, and she thinks the present system serves no one, not even the publishers themselves.

While doing research for an article on digital libraries, Jensen found an article in *Publishers Weekly* with a shocking statistic: 90 percent of the income that a publishing house derives from a book is collected within the first year after publication.

Why, then, does it take at least 75 years for a printed work to officially enter the public domain? (That's a rule of thumb, Hart hastens to remind. In most countries, including the U.S. for works created after January 1, 1978, it's "life of the author plus 50 years" and is therefore completely arbitrary.)

#### "Everyone has this dream idea that electronics is a pot of gold and someday they'll make a mint with their stuff."

And if a book can go for seven decades without making more than a few shekels, then why do publishers have such a strong reaction against the kind of negotiated compulsory license for text and electronic text—a license that might just make them a few more shekels than they would have otherwise—that seems to have worked for other media, such as cable TV?

#### Not quite useless: Copyright Clearance Center

It's not as though the print industry doesn't have an institution in place to deal with licensing rights. It does, but Jensen says the Copyright Clearance Center is only useful in a limited number of fields like the hard sciences, because those are the kinds of publishers who currently belong. Scientists, after all, rely heavily on each other's research, and we're happy they do.

But, Jensen says, a large percentage of all the copyrighted legal material isn't under the CCC umbrella.

"And it's worse in the humanities," she says. "So for the majority of the people and the material, CCC is of little use. Until there's a way to force publishers to join, there's no easy way to pay them electronic rights. You have to track down every copyright holder who isn't a member and get permission from them for each use of work. That's why there's so little going in on conversion to electronic form, except for public domain stuff" like Project Gutenberg.

Not to be redundant, but why? Jensen says the reason is fear and greed. "They're reacting from the gut. They think, 'My stuff is going to get stolen,'" she says. "They don't know the statistics, and even if they did, they wouldn't listen to them. Everyone has this dream idea that electronics is a pot of gold and someday they'll make a mint with their stuff."

She says publishers don't understand that electronics is not a pot of gold. "It in no way increases the pie of information buying dollars out there," Jensen says. "The amount of money out there to buy information is static, just like the rest of the economy. The only thing electronics will do is shift it, or create a marginal increase if the prices are very low."

### It has to be comprehensive

Jensen is convinced that electronic publishing will be stalled until there's some kind of comprehensive licensing scheme. "The difference with electronics is that it's easy to copy and so time consuming and difficult to pay," she says. "Vendors and owners know that, so don't want to grant electronic licenses." To make it easy to pay the system has to be comprehensive and to be comprehensive, it almost has to be compulsory.

Jensen, in fact, compared various copyright schemes in a journal article called "Making Copyright Work in Electronic Publishing Models" in the upcoming Spring and Summer 1992 issue of *Serials Review*.

Her conclusion is that the best model for electronic publishing would be what she calls a "hybrid negotiated compulsory license" where, essentially, copyright holders for text would be free to set their own prices within an agreed-on range, but this negotiation scheme should be backed up by a compulsory license that determines what the rates will be if the parties can't seem to agree.

Denise Caruso 🙎

**(**:))



## Where There's a Will, There's a WAIS

## Public domain system solves many sticky info retrieval problems

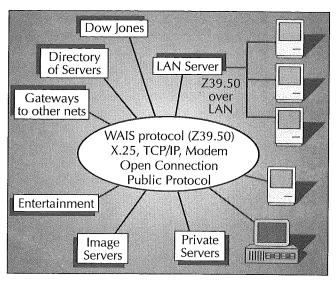
In 1989, Thinking Machines Corp. (Cambridge, MA) — the company that builds the Connection Machine supercomputer — set Brewster Kahle loose on the problem of catalyzing a market for the electronic distribution of information. Finding product (getting people to put the information online) was only part of it. The other part, much more complex, hinged on making sure that customers who wanted product could find it.

So Kahle, who's been with the firm since it was founded in 1983 and is the architect of the CPU of the Connection Machine Model 2, set out to build a system that could navigate the entire panoply of available online data sources, whether on the company's own local area network or on a Unix server halfway around the world.

The project had several goals. First and foremost, the infrastructure had to allow people to make money at electronic publishing. Second, it needed to be elastic enough to support anything from personal computers to consumer electronics devices to supercomputers, at speeds from 1,200 baud to gigabits per second. And third, it needed to be completely accessible and autonomous — *i.e.*, there was to be no single point of control. Anyone who wanted to could snap a server into the network or search a directory of servers for information.

Kahle's project, dubbed WAIS (pronounced "ways"), for Wide Area Information Servers, is already well on its way to achieving those goals.

The electronic library protocol. Even though he says he "changed it completely," Kahle built WAIS around an existing international standard called Z39.50. Once, it simply defined the client-server relationship for a remote bibliographic retrieval system. Kahle's modifications included adapting it for use on global distributed networks and adding multimedia and large document capabilities. (To be searchable, of course, data other than text has to be tagged with text.)



The WAIS Architecture

The modified Z39.50 protocol is now endorsed by the Library of Congress, Apple, Sun Microsystems, Dow Jones and Mead Data Central.

**Unhinging client and server**. What's unique and powerful about the WAIS protocol is that it has unhinged the connection between the client computer's user interface to the system, where a user like me originates a request for information, and the server that translates and acts upon it.

This facilitates a number of important things. As a user, all I have to learn in order to look for information on the network is the WAIS user interface. I don't have to learn to navigate the Internet, or Nexis or Lexis, or CompuServe, or even my own corporate SQL database. If a server is registered on the WAIS system, all I have to do is type — in English — a few words about what I'm looking for in the WAIS text field labeled, "Look for documents about."

The client computer encodes my request in "WAIS-speak" and sends it out to the worldwide network of WAIS servers. The servers translate the query (I don't need to know how), find articles they think match my request, and send them back to me.

I look at the "hits" and select those that look most applicable. Based on that feedback, I can send the query out again until I've gotten the specific information I was looking for. (The method is called "relevance feedback," and it's proven to be a very efficient way to hone in on information.)

**Charging for information.** This "decoupling" of the client from the server is also a powerful tool for electronic

publishers who want to charge for their information in a variety of ways. Control of the charging structure, says Kahle, is resident at the server.

Though how that will evolve isn't clear yet, he says, the technology doesn't stand in the way of any method that publishers might want to use, including subscriptions, document transfer fees, or any combination they can think of. If a piece of information I'm looking for happens to be available for a fee, the server will send a message back to me telling me so and maybe asking me to take some action authorizing the purchase if I haven't already tacitly done so.

Keeping within the law. Another of the system's many useful features (which are far too numerous to list here) is its ability to create "document pointers." These pointers, which stand alone from the document itself, note where a piece of information is located in the network without making a copy of the document itself.

Thus, I can store a pointer so I can return to a document easily, or pass along a pointer to someone who can then find that document. Neither of us is violating copyright laws because we haven't copied or passed along an actual document to someone else. It's the electronic version of the International Standard Book Number, or ISBN, that today makes it possible to locate books in the physical universe.

I can also set up what's called a "dynamic folder" for ongoing topics of research or investigation, whereby WAIS will constantly or periodically (it's up to me) update the folder with new material. All in all, a rather nifty pressure nozzle for the information firehose.

### Gating issues

The reason why Kahle's company would bother going to all this trouble is that it believes the Connection Machine is the very best database server for this nascent electronic publishing market (of course).

But unlike what some companies have and might have done under the same circumstances, Thinking Machines has placed the WAIS source code in the public domain to promote widespread use of the protocol no matter what brand the server.

Kahle realizes the danger of trying to set a standard so early in the market development cycle. But, he says, "I'm working with standards committees to make sure that what we're doing is not a proprietary system, but also that it's free to evolve."

The value of the WAIS system is evident by how quickly new servers are connecting. As of this writing,

145 Internet servers were connected via the WAIS protocol; when we spoke with Kahle less than a month earlier, the number was 120.

A Unix interface already exists, as do VT100, Macintosh and NeXT versions. Dos and Windows are on the way. The Library of Congress, which boasts 25 terabytes of data, has plans to make its catalog available via the protocol.

No "for pay" servers. Despite Thinking Machines' desire to jump-start an electronic publishing market with WAIS, there are still no external interfaces to "for pay" services published in the public domain, though some are under development. Such published interfaces are vital if the WAIS system is to be useful as a real publishing system, not just as a nifty new trick for hackers to play with.

However, progress seems destined to proceed apace, as Kahle and other WAIS supporters gathered at Research Triangle Park in North Carolina on February 3 and 4 to launch the North Carolina WAIS Initiative consortium. The consortium will encourage broad-based development of WAIS interfaces and services and support WAIS freeware.

## "The New York *Times* Review of Servers"?

The potential of the WAIS architecture lends itself directly to the question of what kinds of information might start pouring onto the net. Many established companies, including Dow Jones and Apple, are working closely with Kahle. Dow Jones, in fact, is putting up a test "for pay" WAIS server on its DowVision network with the Wall Street Journal, Barron's and 450 magazines.

Today, published WAIS servers include a directory of servers, the CIA World Factbook, a partial patent database, databases on molecular biology, a poetry server at MIT, cookbooks, descriptions of government software packages, and weather maps and forecasts.

The potential for the number of published WAIS servers to mushroom in number even caused *Byte* magazine to speculate that an independent agency such as *Consumer Reports* might create a rating service to monitor and rate servers in the WAIS directories. The idea would be to serve as an independent guide to quality, publicly tagging servers that regularly delivered bad information or didn't work at all.

What a different world that would be!

Denise Caruso 🙎



## The Reading Room

## A provocative new model promoting open access to information

Brewster Kahle, director of the WAIS project (see story on page 5), offers his views on the clash between private ownership of information, and public access to information.

When written material is distributed over wires, the difference between a public library and a bookstore becomes fuzzy, raising difficult problems for the venerable tradition of free access to information in the public library system.

Public libraries have served two clashing goals: to allow open access to published information and to archive written history for future scholars and posterity. This uneasy combination has been joined into a single institution — public libraries, both centralized and branch — due to the technology of distribution of information on paper.

The wealthy buy. Books and magazines do not have some of the potential problems of electronically distributed material. They are inexpensive and relatively difficult to reproduce, so the wealthy tend to buy copies rather than using the free copies in libraries. Therefore, publishers of most types of material do not incur significant losses in sales because of library copies.

The library's role as archiver is not catastrophically sacrificed by the risk of loaning out books, since the limited public use does not destroy the books. As online distribution (so-called "electronic publishing") becomes common for different segments of written material, access and archiving become quite different, since making a copy is easy and inexpensive, and can be delivered without requiring a person to come into the library.

### A unique way to serve the public

The result of this technology change can be an exciting one where the public library system can refine its charter and serve the public in a widespread way that was not possible with paper. The unique aspects of libraries —

service-oriented staff, lack of profit motive, prevalent locations, and the role in schools — can give them a more important role in the future than they ever had in the paper era.

This essay will suggest a new model for the access goal of the public library, the "Reading Room," that protects and promotes publishers while serving the public in the tradition of the branch library.

Briefly, a Reading Room would offer patrons convenient access to all published information — in printed form or through a screen display — in many places in a town with the help of librarians, but I think the best way to explain the idea is to use the analogy with the U.S. banking system's transition from bank buildings to automatic teller machines:

#### **Old Banks**

Centrally located Limited hours Local information Physical savings books Used rarely

#### Libraries

Centrally located Limited hours Local information Physical library cards

Friendly librarians
Comfortable reading chairs
Current fiction

Easy browsing

#### **Automatic Teller Machines**

Everywhere Available all the time Global information Access cards Constantly used

#### **Reading Rooms**

Everywhere Available all the time Global information Access cards

but still similar —
 Friendly librarians
 chairs Comfortable reading chairs

Current fiction Easy browsing

Reading Rooms can become an active part of people's lives, much as ATMs have transformed people's relationships with banks.

### How the system might work

A minimal design for a Reading Room would be a one-room storefront that had a few comfortable chairs, a bookshelf with current fiction, several computers for browsing and reading, a printer and binder for printing copies on request, and a librarian during regular hours. Interlibrary loan would be used when a particular volume was needed. Therefore, these Reading Rooms would be used by the vast majority of the population, while the existing central libraries would be dominantly used by specialists and archivists. The rest of this essay will outline how this might work.

Benefits of going on line. Today's library patrons, as in the population at large, seem quite willing to use newer technology if it offers a savings of time or other benefits. The widespread preference for online card catalogs has been shown by the large-scale conversion from physical cards despite the costs.

Furthermore, good card catalogs that are accessible remotely have become very popular. The University of California's Melvyl system gets one-third of its requests from people not in the library and often from around the globe, according to Clifford Lynch, director of library automation in the University's Office of the President.

Allowing people access to card catalogs in Reading Rooms can expand the collections that can be accessed, and make them easier to use. People can also use these catalogs from home by dialing up with personal computers or through kiosks in public places.

## Not a replacement for stacks

Catalogs, however, do not replace the stacks for browsing. Computers are making gains in browsing and serendipity — the act of stumbling across something you weren't necessarily looking for — as the screens get bigger and sharper, as full text becomes available, and as better searching systems are developed.

In a 1989 study by Bellcore, researchers found that serendipitous learning of journal information was enhanced by using a computer rather than paper. Experiments in Japan that have used computers to replicate the experience of looking at bookshelves have had encouraging results. We still have a ways to go to improve computer browsing, but great strides are being made.

Printing on demand. When a patron has selected a book or paper to take out of the library, printing and binding can be done on demand. Current screen technology can be used effectively to find useful documents by allowing the patron to browse and read snippets, but output from a printer will still be preferred for reading long pieces.

Although this may seem wasteful, I suggest it is not at all, for if paper were printed only if it were to be read by someone, then there would be many forests left standing. As it is, many pages of books, magazines and newspapers are never even read once.

We can trim paper consumption by encouraging people to print only what they intend to read by previewing on screens, and by printing pieces of long documents, such as chapters. As screens improve, more reading will be comfortable without resorting to paper.

## Royalties and economics

What about the economic issues of royalties? Since the Reading Room monitors all the books printed and viewed, it is easy to keep a record of this and reimburse the publishers. Bear in mind that the current copyright laws are intended to ensure that authors, editors and publishers get a fair stake from the use of their work.

Reading Rooms can easily help in this process. Because they take care of the distribution and retail handling and the printing of the works, publishers can lower their costs of putting out a new volume. This cost savings can cause a flourishing of works intended for a more specialized audience.

**Containing costs.** How about covering the costs of the library? The total cost of maintaining a public library is about \$3–7 per book per year, if all the books in the library are totaled and divided by total overhead. Conservatively, public libraries cost a suburban family of three about \$50 per year.

If we were to give each family in a community a \$50 debit card that could be used in any Reading Room, then every person would have incentive to use that card to read, check out and enjoy published works. If someone needed another card, it could be awarded based on the librarian's assessment of how the last one was used. Since not all families would use their full amount each year, this would guarantee that costs could be contained.

A Reading Room could be very inexpensive to build and operate, assuming a thrifty approach. Two computers, a laser printer, comfortable chairs and bookshelves could be a \$25,000 investment; a well-paid librarian would be another \$50,000 per year; adding \$4,000 for rent would leave the sum well under \$100,000 each year.

**Is it feasible?** Many medium-sized city libraries have a \$3–7 million budget. The same amount of money could fund 30 to 70 Reading Rooms, distributed throughout a city. Obviously these numbers are estimates; my intention is to show that such a plan is feasible.

Might it happen? I don't know. Reassessing our basic institutions in the light of new technology is not something governments do very well.

However, there are many creative people working on the problems of decreasing literacy and falling school performance. Since Reading Rooms make quality information more available to more people, they might help spur learning and exploring among our citizens. I believe the idea is worth further study and possibly a few tests.

Brewster Kahle





#### READERS RESPOND

(continued from page 2) today's network technologies to what we'll need for the commercial reality of digital media. Digital media are being developed on computers, are they not? And computer networks are digital to the desktop. They are also "packetized," switched, two-way and increasingly high-speed. So I will proceed assuming some extension of today's computer networks will best serve digital media.

## The information infrastructure

There are (in round numbers) 100 million computers on earth, about one for every 50 people. Most of these computers are personal and isolated, but suddenly 10 million of them have become connected through an internationally standardized high-speed networking technology called Ethernet.

By "suddenly," I am speaking in geologic time and mean over the last 10 years, and with "high-speed," I am bragging and mean 10 megabits — two books of ASCII text — per second (10 Mbps).

Forerunners for digital media.

The High Performance Computing and National Research and Education Network Act of 1991, signed into law by George Bush last December, aims to invest hundreds of millions of dollars in a high-speed digital backbone for interconnecting computer networks. These, I claim, are the forerunners of future networks for digital media.

This infrastructure of computer networks has been growing since 1969, the dawn of computer networking, when the U.S. Defense Advanced Research Projects Agency (ARPA) began development of computer packet switching technology in the ARPA Computer Network (ARPA-NET). ARPANET has evolved to become a rapidly growing worldwide computer network with more than two million users, many at workstations on high-speed local-area networks (LANs). These LANs are interconnected using lower-speed telephone circuits in a network of networks called the Internet.

Mostly outside the Internet, but increasingly connected to it, are well over a million personal computer LANs. The majority of these are Ethernets, to which 10 million PCs are connected. These Ethernets carry the bulk of traffic in today's information infrastructure.

So, are today's computer networks already a bottleneck?

## Networks are not today's bottleneck

If you read the computer networking trade press, you will get the impression that today's computer networks are a bottleneck. And you will be given the unconvincing assurance that much faster networks are just around the corner.

But beware. Most of what you read flows from network users who mistakenly attribute system speed problems to their networks. And the rest of it flows from network suppliers who find it convenient to use speed to sell higher-margin "future-proof" networking technology.

Watch the silence. However, in the vast majority of cases, today's LANS, especially Ethernet at 10 Mbps, are still nearly empty. The simple



Bob Metcalfe ponders whether digital media will clog future networks.

truth is that when network users sit idled, it's because their file server disks are seeking. Or because their PCs are grinding through poorly implemented or incompatible network protocols. Or their operating systems are switching and swapping. If you doubt this, just hang a scope on the network cable of a bogged-down system, and watch the silence.

#### Peddling speed

In 1976, I had the pleasure of promoting the Ethernet LAN concept when the big argument in data communication markets was whether anyone needed modems at 1,200 bits per second. It was pointed out that computer users could not read terminal screens fast enough to keep up with even a 600-bps modem. Those selling 600 claimed 600 was enough, those with 1,200 said it wasn't, and I was laughed off the stage for talking about networks more than a thousand times faster.

In 1981, I started selling Ethernet products against the LAN market leaders, then Arcnet from Datapoint and Omninet from Corvus. I was careful to push the fact that Ethernet was an official industry standard, but many buyers were happier knowing Ethernet was four times faster than Arcnet and 10 times faster than Omninet. Ethernet entered hypergrowth and soon Omninet and Arcnet fell off in popularity. All the while, as today, LANs everywhere were almost empty.

When is four faster than 16? Five years later, in 1986, when IBM began shipping its IBM Token-Ring LAN, I sympathized as IBM salespeople were forced into ludicrous proofs that their 4 Mbps was somehow faster than Ethernet's 10. Failing that, despite the fact that its 4-Mbps rings were still empty, IBM was then forced to introduce its 16-Mbps variant. And we all soon may be sorry to see IBM try for the third time to have a fast IBM Token-Ring, this time maybe at 64 Mbps.

In 1992, the networking trade press is full of progress reports on Fiber Distributed Data Interchange (FDDI), an emerging computer industry LAN standard and Ethernet's putative successor. FDDI is ten times faster than Ethernet at 100 Mbps. Of course, FDDI has the same two problems that Ethernet had against Arcnet and Omninet: FDDI is still ten times more expensive than Ethernet, and most Ethernets are still empty.

## Bandwidth to burn

FDDI does not impress Ian Leslie here at the Computer Laboratory of the University of Cambridge, England. Leslie and his students have shown me a closet into which Cambridge's new campus network is routed. The closet contains 12 cables, each with three bundles of eight optical fibers. Since Cambridge can now run 500 Mbps on each of these fibers, Ian's closet can handle a total of 144 gigabits per second (Gbps), or 14,000 books per second. I think of this as insurance for when our Ethernets fill with digital media.

It is easy for a genius to predict that Ethernet's 10 Mbps eventually will be too slow for the majority of network-based information systems. But only a fool will say with certainty in what year, or perhaps even in what decade, the development of comput-

ers, their peripherals and demanding digital media applications will make networks the bottleneck, as they were in the 1970s.

## Applications that fill networks

The historical co-evolution of computers, peripherals, applications and networks has some interesting stories. One of them is about laser printing and the invention of LANs.

Fact is, Ethernet was invented to support laser printing. In 1973, at the Xerox Palo Alto Research Center, we felt that we needed Ethernet to run at disk speeds — a few million bits per second — to feed scanned-in, 500dot-per-inch bitmaps to upcoming page-per-second laser printers. Simple arithmetic showed that 1,200-bps modems were not up to the task. And so the experimental Ethernet was attached to such a fire-breathing laser printer at 3 Mbps in 1974. And in 1979, the speed of the standard Ethernet was upped to 10 Mbps simply because it was not much cheaper to run any slower.

A bit of overkill. But nearly 20 years later, among today's millions of laser printers, few run at a page per second, and fewer still at 500 dpi. Further, thanks to page description languages like PostScript, we do not often send scanned-in bitmaps to laser printers, but very much smaller computer-generated document files, with fonts removed. Maybe this is why most Ethernets, most LANs, are empty.

So what do we now think the network requirements are for digital media? This is a question that I am urging that you now think about. To get you started, let me go through one particularly challenging application, perhaps today's equivalent of laser printing in 1973.

## What about 3DTV?

Stewart Lang here at the University of Cambridge is working on three-dimensional television (3DTV). He was already working on building displays for 3DTV, and now I have him worried about getting 3DTV images transmitted to his displays — I have latched on to 3DTV as a benchmark application for future computer networks.

As Lang started my tutorial on the subject, I was interested to learn that most of the "3D" pictures I see are not truly 3D, not even on all those new "3D" workstations we are hearing so much about. Most are two-dimensional projections of 3D objects. As you move your head or walk about, the projections don't change — Mona Lisa's eyes seem to follow you around.

Remember 3D movies? Two slightly different red and green views of the action are updated at 24 frames per second. Your eyes are 6.5 centimeters apart and combine the views in one audience-wide stereoscopic approximation of real life. But even these are not what Lang calls 3DTV.

No special glasses. Lang defines 3DTV as a screen that several people can view in daylight with no special glasses, and see behind objects on the screen stereoscopically when either the people or the objects move around. This is accomplished by giving your eyes different views depending on where they are.

Holograms present 3D by projecting an almost infinite number of views around a room. Lang tells me that a smaller number of views will suffice.

For a 3DTV screen in a small room, with a 120-degree maximum viewing angle, Lang calculates that 100 views will be adequate. So, he asks, what kind of network have you got that can carry 100 TV channels,



one for each view, to a 3DTV? I ask, how many bits are we talking about here? He says that this is a number increasing with time.

### TV is a moving target

Today in the U.S., normal NTSC television is acceptable. This is 644×483 picture elements (pixels) per frame, updated 30 frames per second. Lang estimates 18 bits per NTSC pixel, so I take out my calculator and get 644×483×30×18×100 to get a bit rate of . . . 16.8 Gbps for NTSC 3DTV.

But of course NTSC is already a bit old-fashioned. Soon HDTV will be the norm and this multiplies out to 120 Gbps. After HDTV, people will soon be unsatisfied with image quality anything less than today's 35mm films, which works out to 750 Gbps in 3DTV — five times Leslie's whole closet.

Your first petabit. Here comes what may be your first real sentence above a terabit. How much storage is needed for a two-hour, 35mm 3DTV film? About 5.3856 times 10 to the 15th bits, or 5.3856 petabits. Excuse all the significant digits, but at these exponents every digit counts for a lot.

So I am thinking that if 3DTV takes off (and Lang is not claiming that it will), then maybe networks will indeed be the bottleneck.

## Mitigating factors: the Three Cs

Just as with laser printing, there may be some mitigating factors between these gross bit rate calculations and what actually gets sent.

**Compression.** None of the above numbers accounts for any image compression. Certainly there is much compression to be done. Current

work with compressed broadcast television at 1.544 Mbps (T1 in telephone talk) indicates there is another factor of 10,000 to play with.

Cameras. Lang may be planning to display 3DTV, and I may be planning to transmit it, but where are we going to get 3DTV images from? The building of 3DTV cameras at these resolutions may prove difficult — I won't say impossible. So, maybe the cameras will be one of the bottlenecks.

Computers. Without cameras, 3DTV images may for a long time be computer generated. If so, then it is likely that a concise PostScript-like 3DTV image language will be transmitted instead of the daunting bit streams above.

So you see that 3DTV may turn out exactly like laser printing. A scary projected application drives a technology development, and then turns out later to be much less scary. These mitigations may give us networkers the time we need to keep up. Maybe networks will not be the bottleneck in digital media. But are you worried?

Bob Metcalfe 🙎

## Minnesotans like fun, too

I have to take exception to Peter Dyson's article "Nintendo Takes Aim" (see Vol. 1, No. 8, p. 17). The tone of the article seems to suggest that because Minnesota residents are against using Nintendo machines to play the lottery, they are somehow anti-fun.

Having lived in both California and Minnesota, I have seen two states victimized by the lottery. The lottery is sold to a state on the basis of state revenue and supplemental income to a state's educational programs. In reality, millions of dollars are being sucked out of the state economy and sent to New Jersey, home of the company that runs the games. In addition, California immediately cut the education budget by the amount of the lottery supplements.

Luckily, Minnesota has not adopted the same "anything goes" attitudes about funding for our schools that California has. It costs the state even more to rehabilitate the 5 percent of gamblers who become addicted. Soon you can purchase a lottery chance over a 900 or a 976 number. This is an inherently bad idea too, and it has absolutely nothing to do with Nintendo.

The bottom line is that nobody wins with the lottery. I have not even begun to describe the social evils of the lottery, but I think all of us have seen people buying lottery tickets with food stamps (which is illegal) or people who look homeless spilling a pile of change on the counter to buy into a dream. This was not the result of some governmental slip, it happens to reflect the attitudes of the voters of Minnesota.

Bryan Menell Exact Systems, Inc. St. Paul, MN

I'm no fan of state lotteries myself. Nevertheless, the amount of money wagered each year indicates that gambling is popular with a substantial segment of the population. Making legal, state-run gambling inconvenient — which is what the opponents of Nintendo and 976 betting want to do — just creates a market opportunity for illegal private enterprises. This is the lesson of Prohibition.

— Peter Dyson



## Apple enters consumer market

## Cupertino, it appears, gets the big picture

Back in June 1991, when Apple was making its first moves toward entering the consumer market, we predicted it would merge its superiority in user interface, product design and system-level tools such as QuickTime with the consumer electronics industry's deep understanding of mass markets and miniaturization (see Vol. 1, No. 1).

When its chairman, John Sculley, took the podium at the Winter Consumer Electronics Show last month, it was clear that Apple had indeed realized that digital media and digital technology are the directions that the world is taking. Now Apple must find a way to ride the tide to success.

Best is not enough. Apple is on the same kind of cusp as it was in 1983, just before it introduced the Macintosh. What Apple didn't know in 1983 was that being ahead of your time doesn't earn you sufficient points to justify either an exorbitant price tag or insufficient marketing or focus. We hope that it has now learned that lesson.

If it has not learned, then no matter how fabulous its new consumer products are, it will fail. And it will not have the grace period that was somehow miraculously granted the Macintosh.

## A convergence strategy

Standards for rich media. Apple hopes to take advantage of the convergence of computer and consumer electronics industries by moving to bridge the two markets. To do so, it must lead the charge to establish common "rich media" standards that span the full range from consumer entertainment to hardcore business applications.

If Apple can play a central role in developing these standards (and we see evidence that it is), it can also play a leadership role in both industries. But to take that central role requires cooperation of everyone from IBM to the major Japanese consumer electronics manufacturers.

Leveraging its expertise. Apple has unique expertise, experience, market position and sensitivity that it should be able to leverage into "convergence products." The first generation of Apple consumer products will be low-cost Macs with built-in CD-ROM drives. To be sold through consumer electronics channels, these are promised before the end of the year.

This is a logical step for Apple as a computer company. Macintosh has always been the personal computer best suited for use by nontechnical people. Apple still has the best low-cost hardware/software platform for interactive, graphic-intensive applications. And schools are eating up the new low-cost Macs—vitally important to the home market, since parents like to buy kids what's in the schools.

Apple also has tremendous image and brand recognition, even among consumers who have never used a computer. And Sculley predicts an "incredible explosion" of CD-ROM titles for personal computers this year. This, in turn, should help build

the market for multimedia computers and "players" (equipped to play titles, but not for general-purpose computing tasks).

#### **Enter the PDA**

The next step, promised for early 1993, is bolder—and therefore bound to be far more controversial. As Sculley announced at CES, Apple is at work on families of special-purpose computer-based devices, with special emphasis on small size, easy portability and wireless communications to link them to the rest of the world.

Sculley calls them Personal Digital Assistants, or PDAs. Specific types of PDAs he mentioned include palmtop "executive organizers," players for electronic books and multimedia players.

Apple will develop PDAs in conjunction with other, probably Japanese, firms. They may be marketed both by Apple under its own label and by the other firms under their own names. In fact, Apple recently acknowledged that it is in discussions with consumer giant Sharp, which makes the Sharp Wizard. It is speculated that this is exactly the kind of deal Apple and Sharp are devising.

Initially, these PDAs will be rolled out in the U.S., but Apple clearly believes that the opportunities are global.

These sorts of products will move Apple into the consumer markets, but they will also create some interesting situations. Cooperating with Japanese vendors is key to designing and building miniaturized products, as well as to establishing media standards for the consumer and computer industries. But Apple will certainly end up competing with its partners, in much the same way as Ford competes with Mazda with different versions of its jointly developed cars and trucks.



## ATG and Kaleida: the connection

For some time now, Apple's Advanced Technology Group (ATG) has worked on inventing underlying technology for use in future Apple products. The hard decisions are—and will be—how to bring some of that technology to market, where and how to form partnerships with other companies, and how to generate broad support for the media, data and system software standards that are needed.

A partial solution to the problem was to form a new division, the Advanced Products Group led by Larry Tesler, to aid in technology transfer. Another part of the solution was the decision to form Kaleida, the joint venture that at this point is only between IBM and Apple (although we do not think it is unlikely that other partners and/or owners will join eventually).

What's not clear yet is how much of ATG's technology will be transferred to Kaleida and how much Apple will retain. The objective is most likely to put essential foundation technologies into Kaleida, but to keep some of the implementation technologies for Apple and Apple's joint venture partners.

But it is vital that Apple not short-change Kaleida. Apple recognizes that, unlike Microsoft, it does not have the power to set standards on its own. It cannot succeed if it is technically isolated from the rest of the emerging digital world.

Apple's whole strategy hinges on broad acceptance of the technologies that grow out of what it is doing. This means that Apple can succeed only if Kaleida succeeds. Apple has to put enough technology into Kaleida, and make that technology available broadly and fairly enough, to rally key computer and consumer electronics companies to the cause.

## Compelling products?

The biggest challenge will lie in the specifics of defining, implementing and marketing products that people really want to buy. Will PDAs serve real and unmet needs? Will they have the right combination of design, functionality and price to be compelling? And will Apple be able to market them?

Meeting real needs. The question of demonstrable benefits is the core one. Every successful new consumer electronics product we can think of provided benefits that were immediately obvious and required no real explanation. A product that must be explained is likely dead on arrival.

It's too much to hope that any of the early crop of PDAs might offer the kind of overwhelmingly obvious benefits that fueled VCRs, CDs, desktop copiers, fax machines and cellular phones. But we hope that at least a few of these devices will prove to be really useful for a reasonable number of people.

There are bound to be some duds as well. We are all still groping for what people really want from all this neat technology. Some of our early guesses are almost certain to be wide of the mark.

**Sex appeal.** Apple should be on firmer ground with design and packaging. If it knows anything by now, it should know how to make complex products accessible to nontechnical people and how to package its products in an appealing way.

Marketing, the final frontier. If the benefits of its new devices are obvious, and if they are attractive and well-priced, then Apple should do well. It has excellent brand name recognition for the kind of products that PDAs purport to be. However, if the benefits must be explained at all, it is hard to see how (or why) Apple will succeed.

Judging from the generally poor job that it has done in explaining the benefits of Macintosh, Apple doesn't appear to have any more magic in this regard than the companies that are already serving the consumer market.

## An organizational challenge

The final challenge for Apple will be organizational. Though no announcements were made by press time, it appears that Apple may set up a separate company for its consumer products.

#### A product that must be explained is likely dead on arrival.

In a company with a long history of intra-company, inter-divisional rivalry and wrenching reorganizations, this will not be a trivial undertaking. One of the lessons that Apple is doubtless learning from Kaleida is that it is very difficult to find qualified executives to lead such a division.

So, there are certainly enough risks and enough challenges to this course to keep anyone awake nights. However, we believe that Apple is headed in the right direction. Bridging computers and consumer electronics is Apple's best chance to become a truly *important* company in the coming decade. It is also its best chance to play a significant role in shaping the evolution towards a digital world.

Jonathan Seybold and Denise Caruso 🙎

## Winter CES short on thrills

#### Consumer Electronics Show is the flattest ever, but digital technology is coming around the mountain

According to people who have been attending the Consumer Electronics Show for the past dozen or so years, this year's winter event in Las Vegas was the slowest in memory. But while there was almost nothing new on the show floor, attenders and exhibitors alike were quite aware that a sea of change is sweeping through their business: new ideas, technologies and concepts in consumer products — most of which are digital — will be arriving soon.

Produced by the Electronic Industries Association (EIA), a Washington, DC-based industry trade association, CES does for dealers what Comdex does for computer retailers — gives them the chance to get close to vendors and their latest offerings. Though much of the EIA's membership was opposed to the move, a lackluster market has caused the organization, for the first time in its history, to open the summer CES show in Chicago to the public. The EIA hopes the move will build consumer excitement for new products.

There were really very few products or trends of note at CES, but those that were notable were quite so. Here's a look at the cream of the crop.

## Apple announces the PDA

Probably the biggest news at CES was also a show first: the kickoff speech by John Sculley, chairman and

CEO of Apple Computer, was the first time the head of a computer company gave the keynote speech at the consumer show.

Even if Apple hadn't recently started a consumer products division, Sculley would have been a solid choice. Frank Myers, chief executive of the EIA, claimed that Winter CES had seen a "significant increase" in computer and home office exhibitors, mirroring the market's increase in home office equipment sales.

Sculley, however, probably gave the audience much more than it bargained for. He laid out for the first time Apple's philosophy and strategy for its entry into the consumer electronics business with what he called Personal Digital Assistants, or PDAs.

"Books with batteries." Sculley's PDA vision includes single-purpose computing devices fulfilling such functions as executive organizers, wireless communicators, electronic books and multimedia players. (A detailed look at Apple's consumer strategy begins on page 12.)

Apple's appearance on the scene is likely to be perceived as an ironic turning of the tables by one CES exhibitor. Franklin Electronic Publishers, the New Jersey company that in a previous incarnation was known as Franklin Computer and was sued by Apple for allegedly cloning the Apple II, beat Apple to the punch by more than five years with an entire line of PDAs that it has been selling since the mid-1980s.

Franklin's electronic reference products — the company calls them "books with batteries" — are clearly far more single-purpose than Apple has planned, however. The company sells everything from crossword-puzzle-solving and speaking dictionaries to medical spelling correctors to Big League Baseball Electronic Encyclopedias, at prices ranging from \$40 to \$400.



#### AT&T VideoPhone 2500

The company is profitable, according to a spokeswoman. Its products are sold in more than 30,000 stores and used in more than 5,000 schools across the U.S.

## What's new? DCC and videophones

The only truly new products on the show floor were prototypes of Digital Compact Cassette (DCC) players and recorders and the VideoPhone 2500 from AT&T.

Preparing for battle. AT&T demonstrated for the first time a picture phone that operates over standard phone lines. At a hefty \$1,500 each, the VideoPhone 2500 won't be a runaway hit come the Christmas season, but the purpose was not necessarily to make a million dollars on this product. The goal, which AT&T met, was to be the first on the telecommunications block to have a video phone that works over standard lines.

It has a 3.3-inch LCD monitor and camera lens attached to a standard telephone handset. But at what AT&T claims is two to ten frames per second of color video, the VideoPhone is not exactly a barn burner. In fact,

during a demonstration, it seemed more like a frame every second and a half; and the delay between voice — which worked just as it does on a regular phone — and image was a good (bad, actually) two seconds long. It was quite disconcerting.

However, until recently most people thought it would be impossible to run video and voice communications over a standard phone line, so even limited video is a noteworthy breakthrough.

Rentals may be the dark horse. It's highly unlikely that AT&T will sell more than a couple thousand of these phones. Their extraordinarily high cost, combined with the video's sluggish performance, isn't exactly conducive to big sales.

However, it does seem likely that the company will have a surprising degree of success with its planned overnight rentals. At \$30 per day, they could be a great addition to special events where geographically dispersed groups of people wouldn't otherwise be able to see each other—family reunions, company meetings and the like. When they finally hit the market, the 2500s will be available for rent from AT&T Phone Stores.

**FYI:** more to come. The buzz on the show floor was that AT&T's decision to show the 2500 at CES was a preemptive strike against a coalition that's prepared its own specifications for a video phone. It's said the coalition is made up of five of the seven Regional Bell Operating Companies, which were not yet ready to show the product.

## DCC for days, but where's Sony?

DCC decks were shown in prototype form by half a dozen companies on the exhibit floor. These decks can play analog cassette tapes. They can also record digital audio with a new compression algorithm that removes all musical information outside of the average human's hearing range. The technology, developed by Philips, Tandy and Matsushita, has been licensed to every major consumer electronics manufacturer.

No "golden ears." Even though the compression method loses data, Philips's tests have shown that only half the people who said they could tell the difference between a DCC cassette and a compact disc were able to choose correctly. Statistically, this group of "golden ears" could have guessed and gotten the same results.

Tandy even hosted a concert by jazz pianist Peter Nero, who, after playing some cuts off his new DCC tape, proceeded to "perform" three numbers with his band, pretending to play the piano (hand-synching?) while recorded music played. As became readily apparent, Nero is a great keyboardist, whether he touches the keys or not. The concert left listeners pondering the obvious question: Was it live, or was it Memorex? (Memorex is a division of Tandy.) In all, it was a very effective demonstration.

With the support of the major record labels already in place, DCC machines will be rolling into stores this year.

Conspicuously absent. Sony's absence from the show was conspicuous, especially with DCC all over the floor. Sony was apparently unable to produce working Mini Disc players for demonstration. (Mini Discs are recordable optical discs, about two inches in diameter.) In addition, Sony has had trouble getting the support of the recording industry, which would have to produce enough prerecorded material to make the purchase of a Mini Disc player attractive.



The Big League Baseball Electronic Encyclopedia by Franklin is priced at \$129.95.

The product line was conceived for mobile applications (car stereos, personal disc players, etc.). However, almost from the start, it was seen as a competing product to DCC. While Philips and Sony have agreed to cross-license the technologies (meaning each can manufacture either or both players), the DCC format has received much stronger support from the recording industry.

#### Audio Home Recording Act

DCC's presence underscored the success of an agreement hammered out between the recording industry and the manufacturers of digital home

#### We Have Moved

Digital Media's editorial offices have moved. Please change your databases and personal information managers, personal digital assistants and personal information appliances to reflect the new data listed below:

Digital Media: A Seybold Report 444 De Haro Street, Suite 128 San Francisco, CA 94107 Phone: (415) 575-3775 Fax: (415) 575-3780 recording equipment about home recording rights. (See Vol. 1, No. 3, for a look at the making of the agreement.) The Audio Home Recording Act, when passed, will codify the right for consumers to record music at home.

With the Act meeting little resistance from Congress, the major negotiators of the bill gathered to bask in their success. Though this gathering was proposed to be a discussion of what the Act would mean to those affected by its passage, negotiators at the session downright avoided any real debate of issues raised by the bill.

No mean feat. Certainly the Act is a monumental achievement. Never before have recording artists, producers and equipment manufacturers sat at the same table without threatening to commit mayhem, much less walked from the table with a historical document. Tandy chairman John Roach is largely credited for pulling together the group, which included the Recording Industry Association of America, the National Music Publishers Association and the EIA.

In exchange for explicit protection from litigation regarding home recording, the Act requires that manufacturers install Serial Copy Management System (SCMS) chips on all digital home recording devices to prevent third-generation copies from being produced. They must also pay a tax on digital recording hardware and blank media. That money will then be distributed to the artists and producers who claim to be losing income from home recordings.

There are larger issues. Gary Shapiro, group vice president for the EIA, said that at the request of Congress, the agreement drew a "bright line" around digital audio home recording and nothing else mattered or was affected.

Certainly every industry has its own *modus operandi* when it comes to

licensing material. But in an age of digital media, where there are few precedents to look to for guidance, every new piece of law will be examined very carefully for its effect on others. The motion picture industry, for example, is already scrutinizing this legislation.

They will surface again. While members of Congress have no interest in further confusing the issues before them (especially technology issues), it is naive to believe that this legislation will not significantly affect the larger question of copyright and digital information. Within a few years, the entire issue will surface again for a new group of players.

No one in this session, however, was willing to address this bigger issue publicly. As one observer said, "They don't want to scare Congress before the bill is passed."

However, the mere fact that such an agreement now exists — whether passed into law or not — gives us hope that multimedia producers, the movie and television industries, the recording industry, publishers and anyone else facing questions of intellectual property in the digital world may actually be able to cooperate in creating agreements, not lawsuits, to determine value, rights and responsibilities.

## Offsite, SkyPix demos digital video

Digital video is not far behind digital audio in terms of both technology and products. SkyPix, for example, was demonstrating its digital direct broadcast satellite (DBS) system at a hotel suite off the Vegas strip. Using a proprietary compression algorithm, SkyPix has been able to cram eight channels of digital audio and video onto a satellite transponder meant for a single analog signal.

With a recent infusion of cash and hardware support from Mitsubishi, SkyPix will be operational, with 80 channels of on-demand, pay-perview movies, by April 1992.

Automatic billing. Each home that subscribes to the SkyPix system will be able to select movies via remote control from an onscreen, online selection guide. The user will automatically be billed for the movie, at prices competitive with video rentals. In addition, agreements with many of the major movie houses allow SkyPix to beam a movie to the home unscrambled (for a reasonably higher price) so that the viewer can record it.

Such an agreement is welcomed by movie producers and studios, which today are completely cut out of video rental revenues. They will receive compensation for the "sale" of the movie to SkyPix customers.

The return information, delivered from the customer's home to SkyPix, is carried out through a modem in the satellite decoder via an 800 telephone number. Since each subscriber is individually addressable by the system, it is possible to contact SkyPix, and for SkyPix to contact individual subscribers, over the television. In addition, subscribers will be able to send each other messages over the system.

## Rumors of Gryphon catch CD-I and CDTV off guard

Philips was a major presence at CES with its Compact Disc–Interactive (CD-I) players. In one of the biggest booths on the show floor, Philips Interactive Media of America (PIMA), the organization overseeing the production of most of the CD-I titles, was demonstrating a broad range of titles for the players.



Philips, however, did not own the floor. Commodore produced an advertising campaign that hit during the show, emphasizing the extensibility of the Commodore Dynamic Total Vision (CDTV) system and CDTV's ability to be upgraded to a full Amiga computer and to network with other machines.

Both were probably caught off guard, however, when one of the show dailies ran an article about a new multimedia player being produced by Tandy, code-named Gryphon. While no one at Tandy would talk about such a product, it is very much in line with past statements from both Tandy and Microsoft. Those statements claimed that such a player was in the works, and that it used some variant of Windows as its operating system. More news will follow on this topic, to be sure.

David Baron and Denise Caruso

## FCC allocates spectrum to interactive video, data services

### Commission changes the rules of the game for local license applicants

On January 16, the Federal Communications Commission authorized the allocation of 1 MHz of the public airwaves to "interactive video and data services." Such services would make it possible to interact with one's television set without using a telephone or in any way interfering with the broadcast of the video signal. Two-

way interaction takes place instead over radio waves through a box connected to your TV.

This action came in response to a petition submitted by TV Answer of Reston, VA. The company has developed a technology that would allow viewers of television programming, regardless of its source — cable, satellite, broadcast, etc. — to "talk back" to their televisions (see following story).

Saying no to speculators. To determine who has the right to use the newly allocated radio frequency, the FCC will hold a lottery in 1992. Licenses will be granted to up to 100 of the top markets based on Metropolitan Statistical Areas (MSA) and Rural Service Areas (RSA) data. This is the same procedure that was used when the FCC was granting local cellular phone service.

The cellular phone licenses were quite a valuable commodity, and a lot of financial speculating took place during that lottery process. Allegedly, local licenses were granted and sold on the same day for huge sums of money.

In order to discourage similar speculation over the licenses for local interactive service, and to ensure that licensees actually build out their service areas, the FCC greatly altered the process.

**Evaluating winners only.** First, the FCC will not evaluate all the applicants before the lottery, as it has done in the past. Instead, it will evaluate only the lottery winners, significantly reducing the time necessary to grant the license.

Second, all potential applicants will have to pay a \$1,400 filing fee, up from \$35. In addition, grantees will be required to establish service in 10 percent of their market within the first year, and in 50 percent of the market before they are allowed to sell their licenses.

Another difference from earlier procedures is that the FCC will grant two licenses per market — ensuring, at least initially, that there is market-place competition. One need only look at the local cable licensees to understand why this policy was enacted. (The U.S. Senate just voted to re-regulate the cable television industry in areas without competition to stem the rising costs and poor service often associated with cable.)

In the right place at the right time. TV Answer petitioned the FCC for this ruling, but the ruling does not allocate the bandwidth only to TV Answer. Any company that is interested in using the bandwidth for interactive video and data services can offer its own system to the local operators and programmers. But TV Answer has a big jump on the competition.

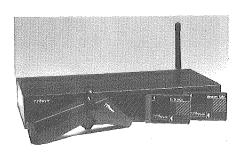
Besides owning and operating the satellite network and processing facilities, TV Answer will be offering all potential licensees assistance in preparing their proposals to the FCC, as well as offering significant technical assistance and financial support for the development of the local cell structure to the eventual winners. The company will also be applying for the local licenses itself in all available markets.

David Baron 2

## TV Answer begs the question

Is it a case of market focus or limited vision?

The first company to propose interactive services over the newly allocated



#### TV Answer system

radio spectrum (see previous story) will be TV Answer. It plans to provide the satellite network and all transaction processing facilities for programmers and local operators of interactive data services.

At the heart of the TV Answer system is a box that sits on top of a television set, housing a microprocessor, a radio antenna and slots for optional memory cards. A remote control device, shaped like a gun with a thumb-operated joystick and a trigger, controls the operations of the TV Answer system (see photo).

How it works. The central hub of the TV Answer network in Reston, VA monitors all programming going out over broadcast satellite (all cable and broadcast network programming is broadcast via satellite before it reaches the home). It is in Reston that the company creates and beams out, in real time, the graphic information that overlays the TV broadcast.

The overlay data is sent by satellite to local operators across the country, who then transmit the signal over radio spectrum to the antenna located on the box (see diagram).

If an interactive option were available, the viewer would see a small icon, such as the capital letter "I," in the bottom corner of the screen. By pulling the trigger on the remote control, the viewer would be given options or questions that could be answered by moving the cursor to the appropriate box and clicking. The information would then travel back

over the airwaves to the local operator, who would beam it back to Reston, where all the transactions are processed.

Advertising overlays. For example, if Chrysler were advertising a new car, it could place a screen over the spot offering two options: "Please have a salesperson call me" or "Please send me more information." Your selection would be sent from your home to Reston, where the data would be collated and sent to Chrysler. In a matter of days, the company claims, you would have the information you had requested.

If it were an advertisement for something smaller, say a CD, there could be a box allowing the viewer to make the purchase immediately. A credit card number and shipping address could be sent to the vendor immediately and transparently to the user. (The user would have to enter a Personal Identification Number to authorize such a transaction.)

## Operating its own communications network

TV Answer plans to operate the communications network, over satellite and radio frequencies, that will allow programmers such as Chrysler or ABC to put interactive information on the viewer screens (over the existing television signal). TV Answer could then process responses from the participants.

The final version of the TV Answer system will contain an Intel 80286 microprocessor (the same chip used in the wildly popular but aging IBM PC/AT). This chip will control all data input and output and will produce all the graphic overlays. There are 16 slots on the front of the box for memory cards, and an optional 20-MB hard drive will be avail-

able. The software is proprietary; no keyboard is necessary or available.

According to the company, the box will be sold in electronics stores by the end of the year at an initial price of between \$300 and \$500. TV Answer predicts it will sell more than a million(!) boxes in the first year, with the street price dropping rapidly as sales increase. A yet-unnamed U.S. manufacturer has committed to producing 1.5 million boxes in the first year of operation.

Limited options. The TV Answer network processes information and transactions in the same way as 800 and 900 number telephone services, only without the telephone. It allows the viewer to respond immediately to advertisements, sweepstakes, public opinion polls, etc., without interrupting the video or reaching for the telephone. And, there is no "telephone bottleneck," which means never having to get a busy signal.

The programming options and services offered by TV Answer will include the following:

- "TV Search" and "Channel Organizer." Both services are provided free to buyers of the TV Answer box. These tools provide an online programming guide and resident information about your particular television options (cable channels, etc.), making search and channel access easier. They also provide simple VCR programming.
- 800-number-type services, including advertisement response, home shopping, political polling, banking services, bill paying, etc. The participant is not charged for these services.
- 900-number-type services, such as sports play-along; play-at-home game shows and sweepstakes; people's choice pageants or award shows; and video games. The user will be charged automatically for participation, but in exchange may

be sent sponsors' coupons or become eligible for prizes and drawings.

As all necessary information about the viewer is stored in the box's memory, there is never the need to give out one's address, phone number, credit card number, etc. By simply clicking on the "Send me more information" button, the user automatically forwards data to the advertiser via TV Answer. Similarly, if you buy something, the amount is automatically charged to your credit card.

Privacy issues. There are currently no guidelines in place to enforce the notion that the consumers must be explicitly informed if the demographic information they provide, willingly or not, is used or sold for marketing purposes. This will probably become a major issue, as two-way, machine-to-machine communication becomes more prevalent. There is already public uproar about two-way services as simple as Caller ID, which simply displays the number of the person calling you.

TV Answer claims to have taken a long look at the issues of privacy and information access that are raised by a two-way link to the home. TV Answer executives say that the company, as managers of the network, will not retain any information (names, programs watched, etc.) that it gathers, but will only facilitate the transfer of information to the programmer.

The program providers determine what questions are asked and what information is exchanged. According to L. Sanders Smith, vice president of video products for TV Answer, the user will be informed explicitly of all information he or she is sending out. However, it is the responsibility of the program provider, not TV Answer, to ensure that this occurs. TV Answer executives are hoping that the organizations and

structures that link directly to the home will police themselves.

TV Question. There is evidence that people want the ability to interact with their television sets, and that they are willing to pay additional fees to get those services. Based on interactive TV trials, consumers seem to like the ability to respond instantly to something on television, or to carry out banking or other services via television.

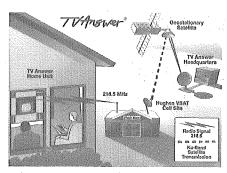
In one year, for example, Canada's Groupe Vidéotron couldn't manufacture its interactive cable TV boxes fast enough to keep up with demand.

But while the Vidéotron system also includes interactive programming, which changes depending on the interests or choices of the viewer, it appears that the services TV Answer wishes to provide consumers are already available through existing equipment. Almost every home has a telephone, and a busy signal and online VCR programming do not seem to be compelling enough reasons for the purchase of a \$300 TV Answer system.

Why not do more? The potential of an inexpensive multifunction PC on every television set would have great implications. But don't talk to the TV Answer folks about this device being a computer — they will have none of it.

To them, "computer" means instant death in the consumer electronics market. They have positioned this system very carefully: as a critical addition to one's home entertainment system, with the capabilities to carry out interactive data operations as well as control all other audio and video equipment you may own. It is not a computer.

It is not a "multimedia" machine either, nor is there any interest in it being compatible with the current crop of CD-based machines that also



TV Answer interactive relay

interface with your television, such as CD-I or CDTV.

All of which seems a shame. TV Answer promises an interactive device, capable of nationwide, instantaneous two way communications, with a powerful processor and a simple interface, but it is limited solely to transactions that can already be accomplished over the telephone. Who would pay \$500, or even \$300, for such a box, and why?

David Baron 🙎

## Correction

The *News* item on page 13 of our January issue concerning a new software video alliance incorrectly listed ICOM Simulations as a participant. ICOM, as we reported in a *Brief* in the December issue, is marketing software-decompressed CD-ROM motion video titles for six hardware platforms. According to its president, Dennis Defensor, however, it is not a member of the software video alliance, and it is not pooling its proprietary technology in any way with any of the companies mentioned.



# Video compression: One step beyond

## International standard wants NTSC quality on a CD-ROM

Last month, we began a discussion of video compression technology by examining the Joint Photographic Experts Group (JPEG) standard for still images. For video production systems where it may be necessary to edit individual frames of a movie, JPEG compression on a frame-by-frame basis is currently the best approach. For most other applications, we can achieve much higher compression by going beyond JPEG. The Moving Pictures Experts Group (MPEG) has proposed an international standard that it hopes will eventually pack NTSC-quality video into the bandwidth available from an ordinary CD-ROM.

As we noted last month, there are really only two tricks for compressing data. First, whenever the data has statistically predictable regularities, it is possible to devise a coding scheme that takes optimal advantage of them.

Examples include run-length coding (sending one example of a repeating pattern plus the number of times to repeat it) and Huffman coding (assigning the shortest codes to the patterns that occur most frequently). Although such codes are perfectly lossless — that is, the reconstructed pattern will be bit-for-bit identical to the original — they can only squeeze the data by a limited amount. For images, compression ratios between 3:1 and 10:1 are about all you can expect.

**Selectivity** is key. Second, in many applications a reasonable facsimile will work as well as an identical

reproduction. By selectively throwing away data during the compression process, an algorithm can achieve much higher compression ratios; reasonable values range from 20:1 to 80:1. The key word here is *selectively*.

The trick is to take advantage of the peculiarities of the human eye and brain and to throw away only the data that won't be missed. The JPEG standard for still images uses some fancy mathematical techniques to tell the difference between essential and nonessential data. It also allows a tradeoff: the amount of loss in image fidelity against the degree of data compression.

Such higher compression ratios exact a cost. It takes lots of number-crunching to compress and decompress the data. Running the computations on today's general-purpose desktop computers takes a long time. However, the impatient (and well-heeled) user can elect to trade cash for time; several vendors offer special-purpose chips and boards to accelerate the computations.

**JPEG video.** JPEG was designed and optimized for compressing still images. Nonetheless, it is possible to apply JPEG compression to each frame of a movie. This makes sense for video producers, who need to save disk space without giving up the ability to edit individual frames. It might also be useful for distributing a few copies of a video — perhaps for beta testing — where turnaround time is more important than the cost of storage media.

In fact, because the viewer is never going to see any single frame by itself, it is safe to apply very high compression and, correspondingly, sacrifice much of the detail. The viewer's eye and brain will be able to reconstruct a satisfying motion picture, and that's all that counts. But even high amounts of JPEG compression — in the neighborhood of 60- or 80-fold — aren't enough to meet the goal of packing NTSC-quality video onto a CD-ROM. For that, we're going to need something on the order of a 200- or 300-fold compression factor.

### Temporal redundancy

In a moving picture, there is usually only a little change from one frame to the next. The background is static and the foreground action takes place in small increments. Many kinds of change occur in mathematically welldefined ways, such as by panning (linear displacement of elements) or zooming.

**Getting the hints.** We can exploit these facts for compression purposes by suppressing some of the frames — say, actually recording only every third or fifth frame — and letting the receiver synthesize (that is, fake) the



## MPEG: the gory details

Three types of framing. The Moving Pictures Experts Group, or MPEG, standard defines three main types of information about the frames of a moving picture. The first is called an Intra-coded frame (I-frame for short). It contains explicit information about every pixel in the image and is quite similar in format to a JPEG-compressed still image. That is, the picture is broken up into tiles that are eight pixels high and eight pixels wide, and each tile is run through a discrete cosine transformation. The results are quantized, then encoded using run-length and Huffman codes. (See Vol. 1, No. 8 for the gory details of JPEG compression.)

However, unlike JPEG, the MPEG standard explicitly calls for adapting the color/detail tradeoff on the fly, tile by tile, based on the content of the image. (As we discussed in last month's Mediascape, JPEG permits, but does not require, on-the-wing adaptation.) This factor alone can improve the compression ratio by 30 percent over a "dumb JPEG" approach.

Because it describes every pixel, an I-frame serves two purposes. First, it erases any accumulated errors that have built up from previous frames. Second, it can be used as the start of a video sequence in interactive applications; a program that jumps around in the movie can begin playing at any I-frame, not just at scene changes.

But by the same token, an I-frame requires a lot of data. For high compression ratios, the objective is to send as few I-frames as possible. The receiver is expected to fill in the gaps between I-frames by interpolation. That is, knowing that a certain pixel was some color in the previous I-frame and is some other color in the current I-frame, the decoder guesses that at each point in time between those frames, the pixel was part way between those colors. Often that is a good guess and the picture doesn't come out too badly. For the parts of a picture where that is the wrong approach, the MPEG standard also provides for two other types of frames.

**Prediction fills in the gaps.** The second type is called a Predictive frame (P-frame for short). It makes no effort to describe the whole picture; it only contains information about changes since the previous I-frame. The picture is broken into tiles  $16 \times 16$  pixels in size, and each tile is compared with the same tile in a previous frame. The system then blithely assumes that

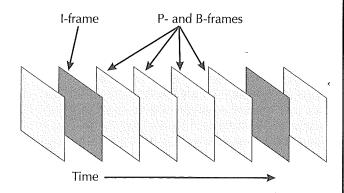
all pixels in a given 16×16 tile are moving together, and it generates a set of numbers that specify the speed and direction of the motion.

That assumption is often true — panning and scrolling are common image shifts — but not always. Thus, a P-frame is really just a "hint" to the decompression system; it reduces, but cannot eliminate, the errors that would be generated by blind interpolation between I-frames.

Back to the future. The third type is a Bidirectional-interpolation or B-frame. Like a P-frame, it contains hints to guide the decoder. However, in using it to build up a displayable image, the decoder will combine its contents with both a past and a future frame. Obviously, the future frame must have been already received, or at least already predicted.

Bidirectional interpolation has the advantage that it deals with areas of the background that are uncovered when a foreground object moves. An area that becomes uncovered can't be predicted from a past frame, but it can be "predicted" from a future frame. Its principal drawback is that it requires buffer memory in the receiver, which has to store at least one earlier and one later version of the displayed image in addition to the version it is currently reconstructing.

Peter Dyson 🙎



An MPEG-compressed movie consists of I-frames interspersed with P-frames and B-frames; the relative frequencies of the various frame types determine how much the movie can be compressed.

#### What about DVI?

When it was introduced to the world by the General Electric/RCA Sarnoff Labs in 1986, Digital Video Interactive offered the only way to obtain good-quality video and sound from a CD-ROM. But despite the marketing efforts of such industry heavyweights as Intel (which acquired the technology in 1987) and IBM, DVI has had an uphill struggle for acceptance.

Now that MPEG's adoption as the international standard seems inevitable, Intel has decided to reposition DVI as a multipurpose video compression engine — the i750 chip set — rather than as a specific encoding technique.

Thus, although the i750 can compress and decompress DVI-coded images, it can also handle JPEG stills and PhotoCD stills. (The next generation of chips will handle MPEG video.) With suitable microcode libraries, it can perform special effects such as merges, fades and wipes.

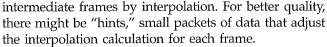
**Two-chip set.** There are two chips in the current product. The 82750PB pixel processor takes a compressed data stream and expands it into a YUV image frame (YUV is a color space comprising measurements for luminance, hue and saturation). It can also compress such a frame into a DVI data stream.

The architecture of this chip resembles a RISC processor — on-chip data registers, pipelined execution, very long instruction word format, multiple loads and stores on every clock cycle — with the addition of specialized function units such as the pixel interpolator to handle zooming and anti-aliasing. A standard microprocessor, says Intel, would have to run at 150 MIPS to match its performance.

The 82750DB display processor takes the YUV frame and massages it for physical display. It handles YUV–RGB conversion; it also contains a digital-to-analog converter and a cursor-shape bitmap. It generates sync pulses and can be programmed for various display resolutions and refresh rates. Current versions of the chip support VGA, 8514/A, XGA and super-VGA resolutions.

These two chips must be supplemented by video RAM and an interface to the host computer to make a complete compressed video product. Intel also supports the chip set by offering software libraries and complete compression-board products.

Peter Dyson 🙎



This technique is a way of trading off the high data rate of raw video for the vast amount of computation needed to create hints and interpolate frames. (No free lunch, remember?) There's another tradeoff, as well. The gain in compression is directly related to the loss of image fidelity.

The limit of what is acceptable will depend on the application. (In other words, you wouldn't use this for medical imaging.) Fortunately, there are many applications where the viewer can be counted on to see what he expects to see and not to notice small discrepancies.

#### NTSC on a CD-ROM

The Moving Pictures Experts Group (MPEG) has sought to achieve NTSC-quality video from a CD-ROM, which would be usable for industrial training films, video games and other commercial information products. There are no MPEG-based commercial products yet, but laboratory demonstrations suggest that its proposed standard will meet that goal.

One reason we began this discussion of video compression with an overview of the JPEG still-image technology is that many of JPEG's techniques are incorporated into the MPEG proposal. Essentially, MPEG combines a variation of JPEG's discrete cosine transform approach (which strips certain small details out of an image) with a set of methods for removing the frame-to-frame temporal redundancy.

JPEG and then some. Ideally, it should be possible to use JPEG-like compression on the first frame of a scene, and thereafter send only information about parts of the image that have changed in subsequent frames. With luck, it, wouldn't be necessary to send another full frame until the entire scene changed — due to a cut to a different camera angle, for instance. The compression ratio could be quite high.

That probably would not work so neatly in practice. Small errors would accumulate, and any loss of synchronization would be disastrous. Moreover, it would make random seeks (e.g., when reading from a CD-ROM in an interactive application) very difficult; the only places you could start viewing would be scene breaks. Both of these factors make it desirable to send a full frame at fairly regular intervals — at least every half-second.

That, in a nutshell, is how MPEG compression works. The data stream contains occasional full frames (com-



pressed using a IPEG-style algorithm) interspersed with tiny packets of hint information. The overall compression depends on what kind of images the movie contains (some compress better than others) and on the frequency of full frames, along with the number of hints.

#### Flexible yet compatible

The MPEG standard is very flexible. It allows the video producer to decide the relative proportion of full frames (which MPEG calls Intra-coded frames; see sidebar for more details) and smaller packets of hints (the Predictive and Bidirectional frames) and thus to control the tradeoff between compression ratios and image artifacts. It supports a wide range of image sizes and frame rates, though it assumes at least 24 frames/second. It allows a wide range of compressed data rates, from the 1.15 megabits per second of a CD-ROM up to the 40 Mbps or so needed for good-quality compressed HDTV.

To assure some interoperability among various MPEG-based products, the standard does define a "constrained parameter bitstream" that all products will support. These parameters include an image size up to 720×576 pixels, a data rate up to 1.86 Mbps and a decoder memory of less than 48 KB.

Like JPEG the standard is very specific about how to decode a signal and says nothing about how to encode one. Thus, there is plenty of room for product differentiation without wild incompatibility.

How about audio? As we go to press, the MPEG standard is nearly complete. The video compression standard is widely agreed on, but there is still some wrangling about how to handle the audio portion of a movie.

This is not a trivial issue. One of the MPEG committee's original goals was to get NTSC-quality video and sound off a CD-ROM, a device whose data rate was originally designed to be sufficient for stereo sound alone. Currently, the best compression algorithms for audio yield only about 4:1 or 5:1 ratios, which leaves only 60 to 80 percent of the bandwidth for the video. It is not clear how this will be worked out.

Until there are commercial products in the field — or at least prototypes in beta testing — we won't know how good MPEG really is. Already, some observers have suggested that the current standard proposal is not optimal at higher data rates and image resolutions. The MPEG committee has already begun a new phase of study focused on digital media in the range of 5-10 Mbps, which may lead to an MPEG-II standard in a year or two.

Peter Dyson 🙎

### And QuickTime, too

QuickTime is widely perceived as the technology. that allows users to play back choppy audio and video in a small window. It is, however, much more complex than that. It is a system-level manager of dynamic data types (data over time, such as audio, video, animations, etc.), hardware peripherals (digitizing boards, VCRs, etc.) and compression algorithms. By putting these tools at the system level, instead of making them particular to each application, QuickTime enables any program, running on any Macintosh with a 68020 processor or better, to take advantage of the new capabilities. In addition, it makes application development significantly simpler and quicker.

The initial release of QuickTime includes four basic components: the system software; a new, dynamic data file format called a "movie"; three Apple data compressors (a JPEG photo compressor for still images, an animation compressor, and a video compressor); and interface standards.

Apple understood well the need for a modular system, one that can easily adapt to the widely varying needs of its customers. The compression manager, for example, allows developers to work with whichever algorithms are appropriate, without worrying about adding individual drivers to the application. In addition to the four software compression algorithms that Apple includes as part of the package, add-in boards and software packages are being developed that support JPEG, MPEG, DVI, P\*64, Group 3 fax and PhotoCD, to name just a few ingredients of the current compression alphabet soup. QuickTime doesn't care which compression scheme you use; the compression manager, will automatically use the algorithm requested by the application, assuming it is available.

The important thing to remember about Quick-Time is that future generations of software will include significantly improved capabilities. Where one might be disappointed in the small window of bad video, before long the Mac will have full-screen (640×480 resolution) video running without additional hardware. In addition, the extensibility of the Macintosh system will provide the groundwork for applications that have not yet even been contemplated.

David Baron 🙎